

INTRODUCTION TO CARRIER ETHERNET VPNS: UNDERSTANDING THE ALTERNATIVES

Table of Contents

Executive Summary	3
Introduction	3
VPN Overview	3
Layer 2 and Layer 3 VPNs	4
VPN Components.....	4
Transport Tunnels.....	4
VPN Labels	4
Pseudowire	4
VPN Services	5
Service Delivery Topologies.....	5
The Multicast Challenge	5
Ethernet VPN Overview	6
Junos Operating System Ethernet VPN Implementations	7
Juniper Networks Point-to-Point VPNs	7
Juniper Networks VPLS Implementations.....	8
Multicast Delivery Using VPLS.....	8
Juniper Networks Products.....	9
Conclusion	9
Appendix A: Ethernet VPN Services and Standards.....	10
Appendix B: CCC and TCC	10
About Juniper Networks.....	11

Table of Figures

Figure 1: VPN overview	3
Figure 2: VPN components	4
Figure 3: Multicast traffic using multiple point-to-point LSPs.....	5
Figure 4: Multicast traffic using point-to-multipoint LSPs.....	6
Figure 5: Ethernet VPN packet overview.....	6
Figure 6: How sites view the VPN.....	8
Figure 7: Multicast traffic using point-to-multipoint LSPs.....	9
Figure 8: CCC/TCC usage	11

Executive Summary

There are a wide range of VPN technologies available. MPLS-based VPNs are the most prevalent technology types used today, with many based on the use of Ethernet transport to provide high-speed communications. This paper describes the various MPLS-based Ethernet VPN services and technologies which Juniper Networks® supports. These include Juniper’s best-in-class VPN implementations, as well as alternatives which provide interoperability with deployed non-Juniper products. This paper is intended for marketing managers seeking to understand Ethernet VPN options, as well as technical managers seeking an overview of their technical alternatives.

Introduction

This paper describes the Ethernet VPN capabilities supported by Juniper Networks. Although Ethernet VPNs can be implemented using various methods such as generic routing encapsulation (GRE), IPsec, Ethernet VLAN stacking, and MAC-in-MAC, the focus here is on the prevalent MPLS-based techniques.

VPN Overview

VPNs partition the resources of a single physical network into multiple logical networks that offer connectivity between different customer sites. Each logical network that links the sites belonging to a customer has a common set of properties such as addressing, services, and traffic forwarding that are private or limited to the scope of that particular logical network.

VPNs consist of three types of nodes, as illustrated in

- Customer Edge (CE): The equipment residing at the customer location. It may be owned and operated by the customer or by the service provider.
- Provider Edge (PE): The equipment at the edge of the service provider “backbone” network. For resiliency, the CE typically connects to one or more PEs.
- Provider (P): The equipment inside the backbone network. This equipment has no awareness of VPNs.

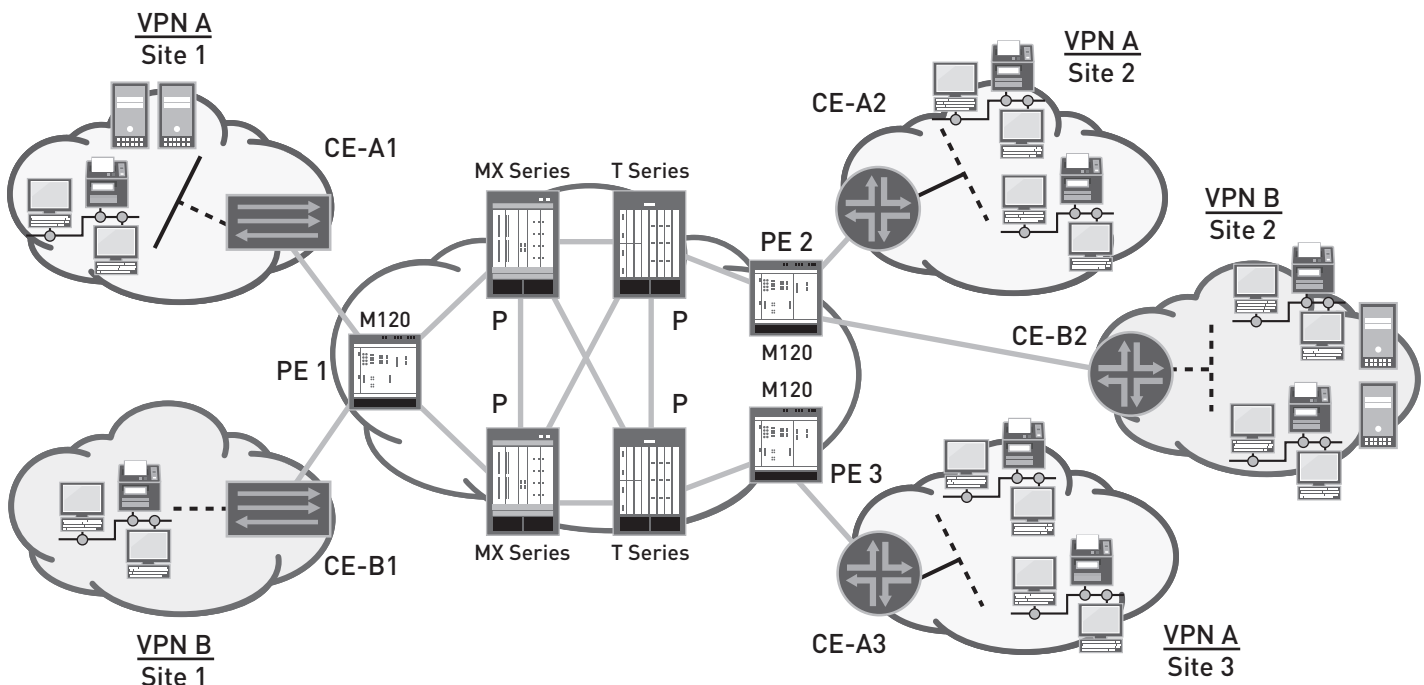


Figure 1: VPN overview

Layer 2 and Layer 3 VPNs

VPNs are classified as either Layer 2 or Layer 3. In the case of a Layer 2 VPN, the provider network offers only transport services between the CEs of the VPN. The routing and peering takes place between CEs; the provider network itself is oblivious to the IP addressing and internal organization of the customer network. This type of VPN is also known as the overlay model. Traditional Layer 2 VPNs include Frame Relay, ATM, or time-division multiplexing (TDM) networks. Modern Layer 2 VPNs use IP/MPLS across the provider network.

In contrast, Layer 3 VPNs have the CEs peering and exchanging routing information with the directly attached PE devices. The provider network can present each customer (or logical network) with route distribution and transport services. Such a model is referred to as the peer model.

Selecting between these models depends on the level of service provider involvement in the customer's network operations. If the customer's goal is to use the provider network only for data transport, a Layer 2 model is better suited since the IP addressing and CE maintenance remains the customer's responsibility. This is more common for large enterprises. The Layer 3 model is appropriate if there is a requirement for the network operator to configure and maintain IP addressing for the customer, which is more typical when the customer is a medium-sized business.

Although Ethernet can be used as the underlying transport mechanism for Layer 3 VPNs, the focus of this paper is on Layer 2 VPNs.

VPN Components

As illustrated in Figure 2, any VPN consists of three major components—the transport tunnel, the VPN, and pseudowires. All three originate and terminate at the PE.




Figure 2: VPN components

Transport Tunnels

A transport tunnel carries traffic between PEs. A separate transport tunnel between each pair of PEs can carry traffic for multiple VPNs, or there may be a separate transport tunnel for each VPN. These transport tunnels are MPLS label switched paths (LSPs). There are two types of LSPs—point to point and point to multipoint. Unicast traffic is delivered using point-to-point LSPs, which provide connectivity between two PEs. Multicast traffic is natively delivered using point-to-multipoint LSPs, which create an efficient distribution tree.

Each LSP transport tunnel is unidirectional. Unicast VPNs use parallel point-to-point transport tunnels to carry traffic in each direction, allowing for bi-directional communication. In contrast, multicast VPNs using point-to-multipoint LSPs do not have a companion carrying traffic upstream.

VPN Labels

The second piece is the VPN label, which specifies the group (enterprise customer, for example) to which this traffic belongs. VPN traffic is carried between PEs using the transport tunnels. A single VPN can deliver both unicast and multicast traffic simultaneously.

Pseudowire

Finally, the pseudowire connection logically connects CEs within the network. A pseudowire begins and terminates at the point (physical port or logical interface such as a VLAN) where traffic enters the PE. There may be several pseudowires carried within each VPN. The IETF PseudoWire Emulation Edge to Edge (PWE3) standards define how Layer 2 traffic is carried across the network. The pseudowire label includes an IETF-defined pseudowire “control word.”

VPN Services

VPN services fall into three categories:

- Line (point-to-point VPN)—for connecting two points, including branch offices connecting to a data center.
- Tree (one-to-many VPN)—typically used for multicast distribution such as IPTV.
- LAN (any-to-any VPN)—for backbone networks connecting major sites. This model is preferred by large enterprises that manage their own IP addresses and are looking for simple connectivity and guaranteed bandwidth between sites.

Service Delivery Topologies

Perhaps the most common VPN model is *hub and spoke*, where numerous locations connect to a central site such as the data center. This is accomplished simply by building a point-to-point VPN (line) from each remote site to the main site, which in turn redirects incoming traffic to its intended destination.

Another common model is a *mesh network*, in which every location has a direct connection to every other location. In this case, the CE (in a Layer 2 VPN) or the PE (in a Layer 3 VPN) forwards each packet directly to its destination. This can be implemented building a full mesh of point-to-point VPNs. Any-to-any (LAN) VPN technologies simplify the provisioning process. In this case, *full mesh* networks have a connection between every pair of endpoints, while *partial mesh* networks do not. One common partial mesh network is a *ring network*.

The Multicast Challenge

Multicast traffic can also be delivered across these network topologies. For example, a spoke site can send the multicast traffic to the hub, which in turn forwards it to all other spokes.

One challenge is how to efficiently distribute multicast traffic. This is most important when there is a large amount of multicast traffic emanating from a central site, such as for delivering IPTV service. There are two ways to deliver this traffic. First, the traffic can be replicated at the ingress point and sent across multiple point-to-point LSPs towards the ultimate destinations. The challenge to this model is that the exact same multicast traffic is carried multiple times across the network, consuming valuable bandwidth. For example, on the six-node metro ring depicted in Figure 3, three copies of each TV channel would traverse the first link.

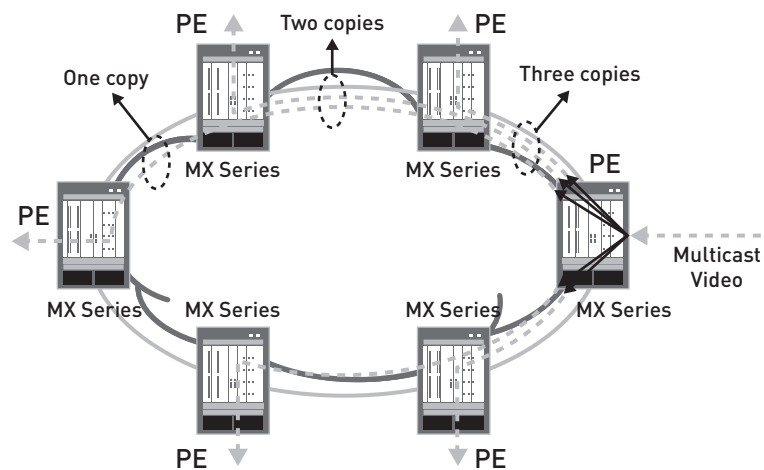


Figure 3: Multicast traffic using multiple point-to-point LSPs

The second approach is to use point-to-multipoint LSPs. In this case, each multicast packet is sent once across the network. Each node delivers the multicast traffic to attached sites, while also forwarding the traffic to the next node. This results in a more efficient traffic distribution, as shown in Figure 4.

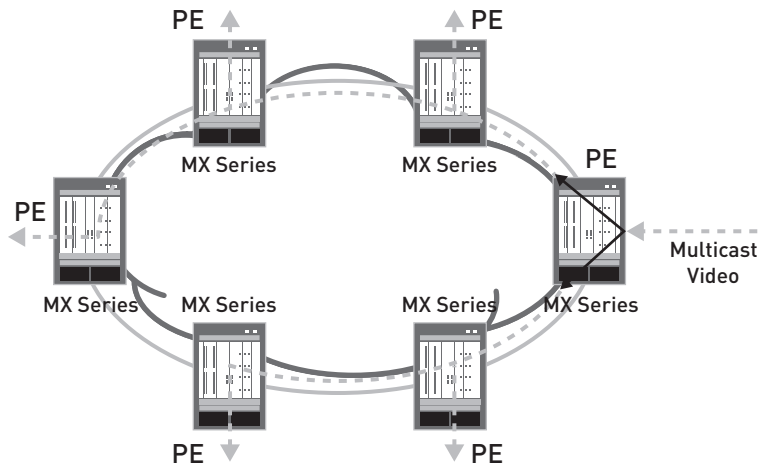


Figure 4: Multicast traffic using point-to-multipoint LSPs

Ethernet VPN Overview

Ethernet VPNs offer end-to-end connectivity between sites belonging to multiple organizations over a shared IP/ MPLS network. This connectivity could be based either on the peer or overlay models. Data coming in to the PE from the CEs belonging to different VPNs is encapsulated with labels for transport over MPLS LSP tunnels. There are two levels of labels that are appended to the VPN data coming into the provider network—an inner VPN label that helps identify the VPN to which the data belongs, and an outer transport label that identifies the outgoing PE to which the data needs to be sent. The outer label is necessary since provider network routers do not maintain VPN service-aware capabilities. These labels are removed before being sent to the CE at the egress end.

Figure 5 shows the packet format of an MPLS-based Ethernet VPN showing fields added by the service provider. If present, there is a different pseudowire control word for each (service provider) 802.1Q VLAN tag. However, many Ethernet VPNs do not include the control word.

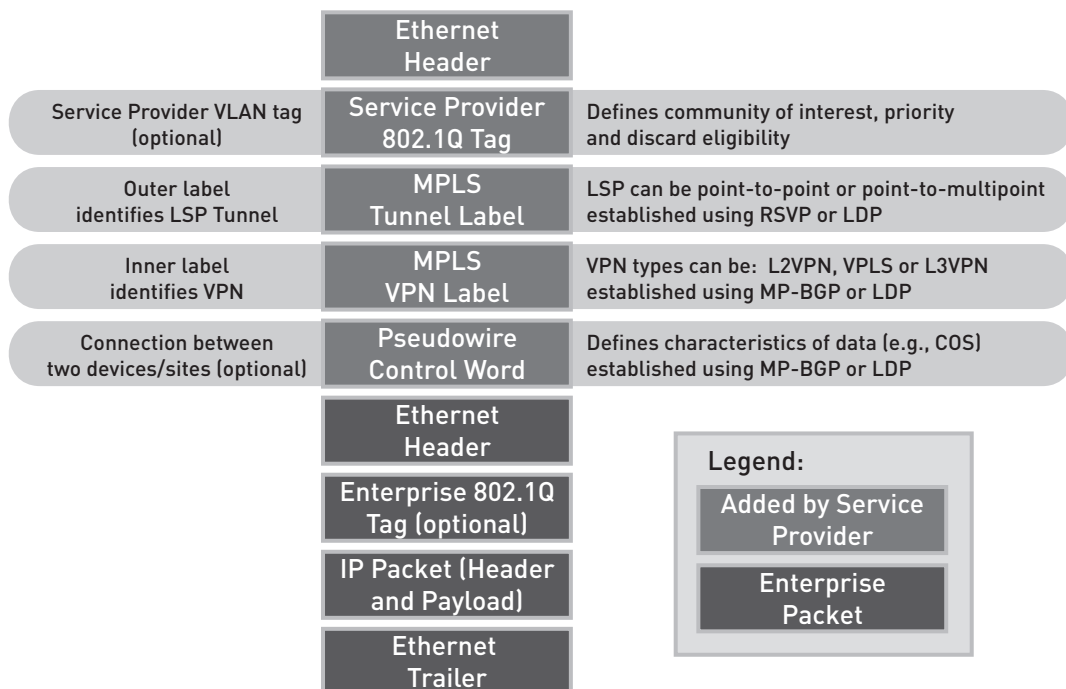


Figure 5: Ethernet VPN packet overview

Junos Operating System Ethernet VPN Implementations

As shown in Table 1, the Layer 2 VPNs supported by Juniper Networks Junos® operating system can be divided into three categories:

- Juniper's recommended VPN family provides superior scaling and a common operational model across VPNs.
- Juniper also supports VPN technologies to ensure the ability to interwork with earlier non-Juniper VPN implementations.
- Finally, Juniper supports legacy implementations which predate any standards but are commonly deployed.

Table 1: Juniper Networks VPN Implementations

DESCRIPTION	LINE (P2P)	TREE (P2MP)	LAN (A2A)
Recommended (RSVP and BGP)	L2vpn	BGP-based virtual private LAN service (VPLS) with point-to-multipoint LSPs	BGP-based VPLS with point-to-point LSPs
Interoperability (LDP)	L2circuit	LDP-based VPLS with H-VPLS	LDP-based VPLS with point-to-point LSPs
Legacy	CCC/TCC		

The recommended VPN family uses RSVP for tunnel establishment, providing reserved bandwidth if desired. BGP is used for VPN establishment (auto-discovery) and for creating pseudowires. The benefits to this approach include:

- Auto-discovery—BGP allows the nodes to signal the VPNs for which they are members and establish the pseudowire connection. This reduces the provisioning steps, making it easier to scale the network and reducing the chance for configuration error.
- Guaranteed bandwidth—Using RSVP allows the operator to offer service-level agreement (SLA) guarantees since bandwidth can be reserved across the network.
- Inter-AS support—BGP supports communication between autonomous systems.

In most other cases, manual provisioning is required and LDP is used to signal the configured information¹.

Juniper Networks Point-to-Point VPNs

Juniper supports three point-to-point Layer 2 VPN implementations—L2vpns, L2circuits, and Circuit Cross-connect (CCC)/translational cross-connect (TCC). These differ primarily in how the layers are established, which in turn affects the capabilities. For the reasons discussed earlier, the BGP-based L2vpn is Juniper's recommended point-to-point Ethernet VPN solution. L2circuits are used primarily for interoperability with non-Juniper equipment. Table 2 summarizes the various Layer 2 point-to-point VPN techniques.

Table 2: Comparison of Point-to-Point VPNs

DESCRIPTION	L2VPN	L2CIRCUIT	CCC/TCC
Tunnel establishment	RSVP or LDP	RSVP or LDP	RSVP
VPN auto-discovery	BGP (dynamic)	N/A (manually provisioned)	N/A (not required)
Pseudowire establishment	BGP	LDP	BGP
COS (EXP bits)	Yes	Yes	Yes
L2 interworking	Yes	Yes	Yes (TCC)
VPNs per tunnel	Multiple	Multiple	One

¹The IETF is considering proposals which implement BGP-based auto-discovery on Ethernet VPNs which otherwise use LDP signaling.

Juniper Networks VPLS Implementations

Juniper supports two VPLS implementations—BGP-based VPLS and LDP-based VPLS. Both are industry standards defined by the IETF. However, Juniper recommends the use of BGP-based VPLS for the reasons discussed earlier.

In a VPLS network, the WAN appears as a LAN to attached devices (CE switches and routers at the edge of the building). All attached locations are on the same IP subnet. Juniper's BGP-based VPLS builds upon the L2vpn implementation described above, allowing the operator to easily offer any-to-any connectivity across the WAN. Using VPLS saves the network operator from provisioning separate point-to-point connections between sites. Instead, each PE port or instance (such as a VLAN) is identified as belonging to a particular VPLS VPN, and the network creates the underlying connectivity. Figure 6 illustrates the network view as seen from one of the VPN A sites.




Figure 6: How sites view the VPN

Juniper Networks supports both BGP-based and LDP-based VPLS implementations, including H-VPLS and BGP-to-LDP interworking. For additional information about VPLS, see the following documents:

- Comprehensive MPLS VPN Solutions, www.juniper.net/us/en/local/pdf/solutionbriefs/3510324-en.pdf
- Virtual Private LAN Service, www.juniper.net/solutions/literature/white_papers/200045.pdf
- LDP-BGP VPLS Interworking, www.juniper.net/us/en/local/pdf/whitepapers/2000282-en.pdf
- Cross-Domain VPLS Deployment Strategies, www.juniper.net/us/en/local/pdf/whitepapers/2000279-en.pdf

Multicast Delivery Using VPLS

VPLS with point-to-multipoint LSPs may be used to efficiently distribute multicast traffic. Most often, this is video traffic being delivered across the backbone network to regional sites—either broadcast TV streams that will be distributed directly to subscribers, or video on demand (VoD) content being sent to distributed caches. This is illustrated in Figure 7. Point-to-multipoint transport tunnels are supported using BGP-based VPLS. For LDP-based networks, H-VPLS is supported.

Alternatively, multicast traffic can be distributed using point-to-point LSPs. This may be appropriate in smaller networks, or if the expectation is that unicast video will predominate.




Figure 7: Multicast traffic using point-to-multipoint LSPs

BGP-based VPLS also allows each site to failover to a backup headend site if the primary site fails. The same VPLS VPN can distribute both multicast (using point-to-multipoint) and unicast (point-to-point) traffic.

For more information about point-to-multipoint LSPs, see Best Practices for Video Transit on an MPLS Backbone at www.juniper.net/us/en/local/pdf/whitepapers/2000106-en.pdf.

Juniper Networks Products

MPLS-based Carrier Ethernet VPNs can be built using Juniper Networks M Series Multiservice Edge Routers, Juniper Networks MX Series 3D Universal Edge Routers, and Juniper Networks T Series Core Routers. These platforms support advanced quality of service (QoS), flexible VLAN tagging and stacking, Internet Group Management Protocol (IGMP) snooping, LAN and WAN PHYs, and point-to-multipoint LSPs. Information about Juniper's Carrier Ethernet support on these platforms is available at www.juniper.net/solutions/service_provider/carrier_ethernet/. While all of these platforms provide carrier-grade networking, the MX Series is designed with Ethernet networking as its primary goal. This platform also supports traditional Ethernet switching including Rapid Spanning Tree Protocol (RSTP), and can be the gateway between Ethernet network segments and the routed backbone. For additional information on Juniper product support for Carrier Ethernet, see Juniper Networks Solution Portfolio for Carrier Ethernet at www.juniper.net/us/en/local/pdf/solutionbriefs/3510342-en.pdf.

Conclusion

Juniper Networks supports a wide range of Layer 2 VPN techniques. The choice of which to deploy depends upon scalability, connectivity, usage, interworking, and interoperability requirements. Point-to-point (line) VPNs can be used to create a virtual network, although large networks can be cumbersome to build. VPLS simplifies the configuration process and efficiently distributes multicast traffic.

Junos OS offers a comprehensive family of simple, scalable Layer 2 VPNs. The recommended solutions use BGP and RSVP signaling, reducing configuration requirements and providing a common operational model across all VPN types. In addition, several additional VPN techniques are available to ensure interoperability with third-party and preexisting VPN implementations.

Appendix A: Ethernet VPN Services and Standards

There are numerous organizations and standards involved in defining VPN services and standards, as illustrated in Table 3.

Table 3: VPN-Related Terminology

STANDARD CONNECTIVITY	RFC 2764	METRO ETHERNET FORUM (MEF)	IETF
Line	Virtual Leased Line (VLL)	Ethernet Virtual Private Line (E-line)	Virtual Private Wire Service (VPWS)
Tree	N/A	Ethernet Virtual Private Tree (E-Tree)	VPLS
Mesh	VPLS	Ethernet Virtual Private LAN (E-LAN)	VPLS

The Internet Engineering Task Force (IETF) defines VPN standards. Information about the various technologies can be found at the following sites:

- MPLS technology including point-to-point and point-to-multipoint LSPs: www.ietf.org/html.charters/mpls-charter.html
- Layer 2 VPNs including VPLS and VPWS (point-to-point VPNs): www.ietf.org/html.charters/l2vpn-charter.html
- Layer 3 VPNs: www.ietf.org/html.charters/l3vpn-charter.html
- PWE3: www.ietf.org/html.charters/pwe3-charter.html

Appendix B: CCC and TCC

Circuit Cross-Connect (CCC) is an early but widely deployed technique which, as the name implies, simply cross-connects two logical interfaces within a router. Although there is no concept of a VPN label, CCC can be used to provide VPN service. Since there is no VPN label, each transport tunnel supports a single customer.

As shown in Figure 8, CCC can provide an MPLS tunnel (top) between sites, similar to L2vpn and L2circuits. This is called MPLS tunneling (top). It can also be used to interconnect different LSPs that belong to different traffic engineering domains (LSP stitching, middle). Finally, CCC can provide local switching within a router without using MPLS (L2 switching, bottom). When using CCC, the ingress and egress ports must use the same Layer 2 protocol such as Ethernet, Frame Relay, or ATM.

Translational Cross-Connect (TCC) is similar to CCC except that it offers L2 interworking support, allowing different Layer 2 media to be connected at each end of the circuit.

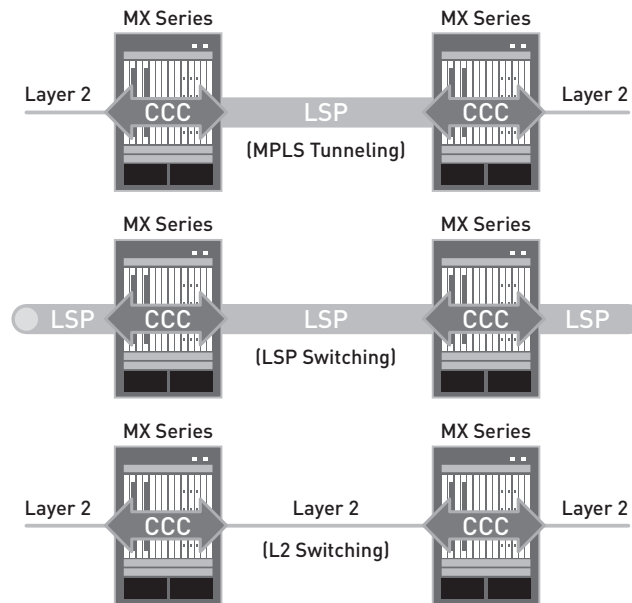


Figure 8: CCC/TCC usage

About Juniper Networks

Juniper Networks, Inc. is the leader in high-performance networking. Juniper offers a high-performance network infrastructure that creates a responsive and trusted environment for accelerating the deployment of services and applications over a single network. This fuels high-performance businesses. Additional information can be found at www.juniper.net.

Corporate and Sales Headquarters

Juniper Networks, Inc.
 1194 North Mathilda Avenue
 Sunnyvale, CA 94089 USA
 Phone: 888.JUNIPER [888.586.4737]
 or 408.745.2000
 Fax: 408.745.2100
www.juniper.net

APAC Headquarters

Juniper Networks (Hong Kong)
 26/F, Cityplaza One
 1111 King's Road
 Taikoo Shing, Hong Kong
 Phone: 852.2332.3636
 Fax: 852.2574.7803

EMEA Headquarters

Juniper Networks Ireland
 Airside Business Park
 Swords, County Dublin, Ireland
 Phone: 35.31.8903.600
 EMEA Sales: 00800.4586.4737
 Fax: 35.31.8903.601

To purchase Juniper Networks solutions, please contact your Juniper Networks representative at 1-866-298-6428 or authorized reseller.

Copyright 2010 Juniper Networks, Inc. All rights reserved. Juniper Networks, the Juniper Networks logo, Junos, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners. Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.